

**METHOD AND APPARATUS FOR IMPLEMENTING POINT-TO-  
MULTIPOINT COMMUNICATIONS OVER A CONNECTION-  
BASED DATA COMMUNICATION NETWORK**

5    **Technical Field**

[0001]        The invention relates to data communication networks. The invention relates specifically to the provision of point-to-multipoint communications in computer networks in which a number of network segments are linked by a connection-based network. The connection-based network may comprise a cell relay network such as asynchronous transfer mode (ATM) networks.

**Background**

[0002]        Computer networks carry data between various devices. The data may be carried in connection-based links, such as the virtual circuits in an Asynchronous Transfer Mode (ATM) network. Data may also be carried between devices in network segments wherein data is broadcast to all of a plurality of devices on the segment on a broadcast-type medium. An example of the latter is an ethernet network. It is typically convenient to set up local area networks (LANs) using a broadcast type medium over which devices can share data.

[0003]        In some circumstances, for example, where a LAN is required to connect devices which are geographically distant from one another, the LAN may be broken into separate segments. Within each segment devices can exchange data by way of a broadcast-type medium. The segments may be connected to one another by way of connection-

- 2 -

based links. Such a LAN may be called a virtual LAN (VLAN). There are various existing standards which apply to VLANs of specific types.

[0004] LANE, described in ATM Forum Specification

5 af-lane-0021.00, describes a method for LAN emulation over ATM.

LANE specifies a client-server architecture in which clients only forward frames with known destinations. All other traffic is forwarded to a server which determines destinations to which the frames should be forwarded.

10 [0005] Various standards exist for transporting variable-sized data packets, such as ethernet frames over ATM connections. For example, IETF RFC 2684 provides protocols for encapsulating various packet types (ethernet, IP, etc.) for transport over AAL5.

15 [0006] Various standards for routing variable-sized data packets over connection-based networks also exist. These include classical IP and ARP over ATM (IETF RFC 2225), multiprotocol over ATM ("MPOA") as specified in IETF RFC 2684 or multiprotocol label switching ("MPLS") as specified in IETF RFCs 2917 and 3035.

20

[0007] Figure 1 shows an example of a simple VLAN 10. VLAN 10 comprises three segments 12 (individually labeled 12A, 12B, and 12C).

Each segment 12 comprises a number of devices 13 connected by an ethernet. Segments 12 are connected by virtual circuits through an ATM

25 network 14. Bridges 16A, 16B and 16C (collectively bridges 16)

interface segments 12 to ATM network 14. Network 14 is typically set

- 3 -

up to provide a point-to-point (P2P) virtual circuit between each bridge 16 and other bridges 16 which belong to VLAN 10.

[0008] The topology of a VLAN may be established using a  
5 protocol such as the spanning tree protocol (STP). STP generates a network topology which is defined by a spanning tree. The spanning tree defines a topology which does not include any closed loops.

[0009] There are various situations in which it is desirable to direct  
10 communications which originate at one point in a VLAN 10 to multiple destinations in multiple network segments. For example, it may be desirable to simultaneously deliver streaming video data from one device 13 to a number of other devices 13 on the network.

15 [0010] There are ATM standards which permit the configuration of point-to-multipoint (P2MP) ATM channels. It has not been practical to use such channels for multicasting data which originates in a segment of a VLAN.

20 [0011] U.S. patent No. 6,111,880 discloses a switch which supports both ATM and ethernet operation. The switch uses a special internal data format. The switch enables multicast data traffic by setting a "multicast mask" tag in a VPI/VCID lookup table in memory.

25 [0012] U.S. patent No. 5,852,606 discloses a method and apparatus for transmitting cells across an ATM switch bus. A routing tag provides multicast group destination information.

- 4 -

[0013] There exists a need for cost effective methods and apparatus for carrying multicast data transmissions which originate at a segment of a bridged LAN.

5

Summary of the Invention

[0014] This invention provides methods and apparatus for delivering multicast data traffic to a plurality of destinations. One aspect of the invention provides a method for delivering multicast data traffic  
10 originating in a broadcast-based computer network to a plurality of destinations on a connection-based network. The method comprises providing a bridge connecting the broadcast-based network and the connection-based network. The bridge provides one or more ports at which virtual channels in the connection-based network can terminate.  
15 The method also includes setting up a point-to-multipoint virtual channel in the connection-based network. The point-to-multipoint virtual channel has a root at a first one of the ports and a plurality of leaves at destination nodes in the connection-based network. The first one of the ports is associated with one or more multicast addresses in a filtering database  
20 associated with the bridge. The method involves, at the bridge, forwarding multicast data frames addressed to the multicast address and originating in the broadcast-based computer network to the first one of the ports. The connection-based network may comprise an ATM network. In some embodiments the connection-based network is internal  
25 to a bridging device and may comprise a switching fabric within the bridging device. In other embodiments the connection-based network is a geographically extended network.

- 5 -

[0015] Another aspect of the invention provides a method for carrying multicast data traffic originating at a source segment of a virtual network to a plurality of destination segments of the virtual network. The  
5 the source and plurality of destination segments each connect to a connection-based network at a bridge. The method comprises: at a first bridge connected to the source segment, associating at least one multicast address with a first remote interface port and configuring the remote interface port as an ingress-only port; provisioning in the connection-  
10 based network a point-to-multipoint virtual channel having a root endpoint at the remote interface port and a plurality of leaf nodes; directing multicast data addressed to the at least one multicast address to the first remote interface bridge port; and, passing the multicast data to the destination segments by way of the point-to-multipoint virtual  
15 channel.

[0016] Yet another aspect of the invention provides a bridge device comprising: a network interface configured to receive variable sized data frames from a first network; a plurality of bridge ports; a switching fabric  
20 configurable to provide data connections between the bridge ports and a plurality of external data connections, the external data connections each associated with one of one or more output interfaces; a point-to-multipoint virtual channel configured in the switching fabric, the point-to-multipoint virtual channel having a root at a first one of the bridge  
25 ports and a plurality of leaves, the leaves each connected to one of the external data connections; and, a filtering database associated with the

- 6 -

bridge, the filtering database containing a first entry associating one or more multicast addresses with the first one of the bridge ports.

[0017] A further aspect of the invention provides a virtual local  
5 area network comprising: a plurality of segments interconnected by a connection-based network; a bridge associated with each of the segments each bridge connecting a corresponding one of the segments to the connection-based network; a first bridge associated with a first one of the segments, the first bridge comprising a plurality of bridge ports each  
10 capable of being connected to a virtual channel in the connection-based network; a point-to-multipoint virtual channel in the connection-based network, the point-to-multipoint virtual channel having a root node associated with a first one of the bridge ports and a plurality of leaf nodes, each of the leaf nodes connected to one of the bridges  
15 corresponding to another one of the segments; and a point-to-point virtual channel in the connection-based network, the point-to-point virtual channel connecting a second one of the bridge ports to one of the bridges corresponding to another one of the segments.

20 [0018] Further aspects of the invention and features of specific embodiments of the invention are described below.

#### Brief Description of the Drawings

25 [0019] In drawings which illustrate non-limiting embodiments of the invention,

- 7 -

Figure 1 is a schematic diagram of a network having a number of ethernet segments connected by cell relay channels;

Figure 2 is a block diagram of a bridge according to an embodiment of the invention;

5        Figure 3 is a schematic diagram illustrating a virtual LAN in which multicast data is transmitted to several destinations; and,

Figure 4 is a schematic diagram illustrating a virtual LAN having a configuration different from that of Figure 3 in which multicast data is transmitted to several destinations.

10

#### Description

[0020]        Throughout the following description, specific details are set forth in order to provide a more thorough understanding of the invention. However, the invention may be practiced without these  
15        particulars. In other instances, well known elements have not been shown or described in detail to avoid unnecessarily obscuring the invention. Accordingly, the specification and drawings are to be regarded in an illustrative, rather than a restrictive, sense.

20        [0021]        This invention provides a bridge 20 as shown in Figure 2. Bridge 20 comprises a local interface port 22, and a plurality of remote interface ports 24. From the point of view of bridge 20, local interface port 22 and remote interface ports 24 may be equivalent. Each remote interface port 24 can serve as a termination endpoint for one or more  
25        virtual channels in a connection-based network. Remote interface ports 24 could, but do not need to, comprise separate physical devices. Remote interface ports 24 may comprise virtual remote interface ports 24.

- 8 -

[0022] In the illustrated embodiment, bridge 20 is part of a device 21 which includes a switching fabric 26 and one or more output interfaces 28. Switching fabric 26 selectively establishes virtual channels (or "connections") between ports 24 and the output interfaces 28.

[0023] Bridge 20 has access to a filtering database 23. Bridge 20 looks up in filtering database 23 the destinations of frames received at local interface port 22. If database 23 indicates that the destination for a frame is on segment 11 then bridge 20 may drop the frame. If filtering database 23 indicates that the destination for the frame is associated with a specific port 24 then bridge 20 forwards the frame to the specific port 24. If filtering database 23 has no record of the destination for the frame then bridge 20 may forward the frame to multiple ports 24.

[0024] In the illustrated embodiment, network interface 22 connects to a broadcast type medium, such as an ethernet network. Data is transferred in the broadcast-type medium in variable-size frames. In the illustrated embodiment, switching fabric 26 is an ATM switching fabric and data is transferred out of ports 24 and is received at ports 24 in fixed-size cells. The frames of network segment 11 may be transported in fixed-size cells using any suitable protocol. ATM adaptation layer 5 (AAL5) specifies one protocol for encapsulating variable length data frames for delivery over an ATM cell relay network.

[0025] As shown in Figure 3, device 21 can be connected between a network which operates in a broadcast-type medium and a connection-



- 9 -

based network 14. Device 21 can be used to establish point-to-multipoint communications between a device connected to local interface port 22 and a plurality of endpoints in the connection-based network 14. In Figure 3 the data to be delivered by way of the point-to-multipoint virtual channel originates in a segment 11 of a VLAN which is connected to local interface port 22 and the endpoints are shown as being in segments 12A, 12B, and 12C of the VLAN. Segments 12A, 12B and 12C of the VLAN are each connected to connection-based network 14 by way of a bridge 16. Bridges 16 could comprise devices 21 or other bridges.

10

[0026] Bridge 20 directs data to be multicast to a port 24A. Port 24A is configured to be at a root endpoint 35 of an ATM P2MP virtual channel 30 which has leaf endpoints (36A, 36B and 36C) at each of bridges 16. Port 24A may be termed a "first remote interface port". In the

15 illustrated embodiment, P2MP virtual channel 30 is provisioned within switching fabric 26 and has leaf endpoints at three different interfaces 28. Virtual channels 30A, 30B and 30C connect interfaces 28 to bridges 16A, 16B, and 16C respectively. Virtual channels 30A, 30B and 30C may be considered to be part of P2MP virtual channel 30 even though they

20 may be provisioned separately. The endpoints of virtual channels 30A, 30B and 30C at bridges 16 may be considered to be leaf endpoints of P2MP virtual channel 30. Virtual channel 30 may be provisioned in any suitable way, either administratively or via ATM signaling.

- 10 -

[0027] The embodiment of Figure 3 can reduce bandwidth required at bridge 20 since it is not necessary for bridge 20 to replicate all multicast frames.

5 [0028] In some cases there is a need for signals to be transmitted from the destination of a multicast to the source. As shown in Figure 3, a plurality of point-to-point virtual channels 34 may be provided for this purpose. Virtual channels 34A, 34B and 34C extend respectively between ports 24B, 24C and 24D at bridge 20 and bridges 16A, 16B, and  
10 16C. Virtual channels 34A, 34B and 34C may carry bidirectional return control messaging. Virtual channels 34A, 34B and 34C terminate at each end at a bi-directional endpoint 37. Ports 24B, 24C and 24D may be termed "second remote interface ports".

15 [0029] Filtering database 23 includes an entry which identifies port 24A as the port to which bridge 20 should forward all frames carrying specified multicast addresses which indicate that the frames should be delivered by way of virtual channel 30. When such frames are received at local interface port 22, bridge 20 forwards them to port 24A and, from  
20 there, carried by P2MP virtual channel 30 to each of bridges 16. Unicast frames can be carried to their destinations by way of one of virtual channels 34.

[0030] Multicast frames which have multicast addresses which do  
25 not correspond to the specified multicast addresses can also be delivered by way of virtual channels 34. This permits differentiated service to be provided to different classes of multicast traffic. For example, virtual

- 11 -

channel 30 may comprise a high bandwidth, low jitter connection. Specified multicast traffic, such as a video broadcast, audio broadcast, multimedia transmission, or the like could be carried on virtual channel 30. Other traffic, which could include other multicast traffic, may be  
5 carried on virtual channels 34. In general the Quality of Service (QoS) provided by channels 30 and 34 may be different.

[0031] Normally when a bridge receives a frame addressed to a multicast address it forwards the frame to multiple bridge ports. In  
10 preferred embodiments of this invention, filtering database 23 includes entries which associate specific multicast addresses with specific remote interface ports 24. For example, filtering database 23 may include entries which associate one or more multicast addresses with port 24. Where such an entry exists for a multicast address then bridge 20 directs  
15 frames addressed to the multicast address only to the specified port. These entries may be configured manually or automatically using a suitable protocol such as GMRP (GARP Multicast Registration Protocol).

20 [0032] It is also desirable, and in some cases essential, to eliminate loops along which data can propagate. One way to do this is to designate virtual channel 30 as a unidirectional virtual channel. The port 24A at the root 35 of P2MP virtual channel 30 may be designated "ingress only". In this context, "ingress" means a direction of data flow from bridge 20 into  
25 virtual channel 30 for delivery to the leaf endpoints 36 of P2MP virtual channel 30. Port 24A drops any data which is egressing from virtual channel 30. In this context, "egress" refers to a direction of data flow in

- 12 -

which data arrives at bridge 20 from virtual channel 30. Ports in bridges 16 at the leaf endpoints 36 of P2MP virtual channel 30 may be designated "egress only". Those ports only pass data arriving at bridges 16 through virtual channel 30 and block data originating at bridges 16 from being sent out on virtual channel 30. Bridges 16 may drop all data which would otherwise be sent out such ports.

[0033] Another way to avoid loops is to either not provide or to disable virtual circuit-to-virtual circuit (VC-to-VC) forwarding at port 24A as well as at the leaf endpoints of virtual channel 30.

[0034] Where virtual channel 30 is part of a network which is configured by a spanning tree protocol (STP) then steps may be taken to prevent the spanning tree protocol from detecting multiple paths between bridge 20 and one or more of bridges 16. STP attempts to eliminate multiple paths. One way to avoid such problems in cases where STP is being used is to effectively disable STP for virtual channel 30. This may be done by configuring STP to ignore any ports which are configured as "egress only" or "ingress only".

20

[0035] Figure 4 illustrates an alternative embodiment of the invention wherein a P2MP virtual channel 30' splits into multiple branches 30A', 30B' and 30C' at a node in connection-oriented network 14 which is separate from device 21. This conserves bandwidth in portions of network 14.

25

- 13 -

[0036] Some specific embodiments of the invention provide a bridge connected to a source of multicast traffic. A P2MP virtual circuit root endpoint is configured on a first ATM bridge port of the bridge. The first port is configured as an ingress only port. A filtering database entry  
5 associates at least one multicast address with the first port. The source bridge may have other ports configured as normal ports.

[0037] Bridges at destination (leaf endpoints) of the P2MP virtual channel have ports connected to the leaf endpoints. These ports are  
10 configured as "egress-only" ports. Bridges 16 are configured to not learn source addresses from data received by way of an egress-only port. The destination bridges may have other ports configured as normal ports. When P2MP virtual channel 30 is set up, the port at which the root of virtual channel 30 is located may be configured automatically to be  
15 ingress-only ports. The ports at the leaf endpoints of virtual connection 30 may be configured automatically to be egress-only ports.

[0038] Those skilled in the art will appreciate that the foregoing embodiments of the invention provide a mechanism for delivering  
20 multicast data traffic over P2MP virtual channels. The virtual channels may be ATM virtual circuits. Bidirectional connectivity is provided by additional point-to-point virtual channels. No ATM VC-merge function is required anywhere in the multicast path.

25 [0039] Certain implementations of the invention comprise computer processors which execute software instructions which cause the processors to perform a method of the invention. The invention may also

- 14 -

be provided in the form of a program product. The program product may comprise any medium which carries a set of computer-readable signals comprising instructions which, when executed by a computer processor, cause the data processor to execute a method of the invention. The  
5 program product may be in any of a wide variety of forms. The program product may comprise, for example, physical media such as magnetic data storage media including floppy diskettes, hard disk drives, optical data storage media including CD ROMs, DVDs, electronic data storage media including ROMs, flash RAM, or the like or transmission-type  
10 media such as digital or analog communication links.

[0040] Where a component (e.g. a software module, processor, assembly, device, circuit, etc.) is referred to above, unless otherwise indicated, reference to that component (including a reference to a  
15 "means") should be interpreted as including as equivalents of that component any component which performs the function of the described component (i.e., that is functionally equivalent), including components which are not structurally equivalent to the disclosed structure which performs the function in the illustrated exemplary embodiments of the  
20 invention.

[0041] As will be apparent to those skilled in the art in the light of the foregoing disclosure, many alterations and modifications are possible in the practice of this invention without departing from the spirit or scope  
25 thereof. For example:

- The invention could be embodied in a LANE-type virtual network having unidirectional connections to LAN Emulation Servers

- 15 -

(LES) and LAN Emulation Clients (LECs). In such embodiments the servers may specify separate paths for multicast source to destination communication and destination to source communication as described above;

- 5 • Output interfaces 28 could comprise mechanisms for adapting a cell stream received on P2MP virtual channel 30 into one or more other data formats such as ethernet frames, frame relay frames, packet over SONET frames etc. In such embodiments, P2MP virtual channel 30 may exist completely within a bridging device

10 21.

- A single bridge 20 may have a plurality of remote interface ports connected to root endpoints of a plurality of P2MP virtual channels 30. Each of the ports connected to one of the plurality of P2MP virtual channels 30 may be associated with one or more
- 15 multicast addresses by way of static entries in a filtering database 23.

Accordingly, the scope of the invention is to be construed in accordance with the substance defined by the following claims.